

NASA/CR-2000-210126
ICASE Report No. 2000-27



Robust Multigrid Algorithms for the Incompressible Navier-Stokes Equations

*Ruben S. Montero and Ignacio M. Llorente
Universidad Complutense, Madrid, Spain*

*Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA*

Operated by Universities Space Research Association



National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

Prepared for Langley Research Center
under Contract NAS1-97046

May 2000

DTIC QUALITY INSPECTED 1

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

20000523 042

ROBUST MULTIGRID ALGORITHMS FOR THE INCOMPRESSIBLE NAVIER-STOKES EQUATIONS*

RUBEN S. MONTERO[†] AND IGNACIO M. LLORENTE[‡]

Abstract. Anisotropies occur naturally in CFD where the simulation of small scale physical phenomena, such as boundary layers at high Reynolds numbers, causes the grid to be highly stretched leading to a slow down in convergence of multigrid methods. Several approaches aimed at making multigrid a robust solver have been proposed and analyzed in literature using the scalar diffusion equation. However, they have been rarely applied to solving more complicated models, like the incompressible Navier-Stokes equations. This paper contains the first published numerical results of the behavior of two popular robust multigrid approaches (alternating-plane smoothers combined with standard coarsening and plane implicit smoothers combined with semi-coarsening) for solving the 3-D incompressible Navier-Stokes equations in the simulation of the driven cavity and a boundary layer over a flat plate on a stretched grid. The discrete operator is obtained using a staggered-grid arrangement of variables with a finite volume technique and second-order accuracy is achieved using defect correction within the multigrid cycle. Grid size, grid stretching and Reynolds number are the factors considered in evaluating the robustness of the multigrid methods. Both approaches yield large increases in convergence rates over cell-implicit smoothers on stretched grids. The combination of plane implicit smoothers and semi-coarsening was found to be fully robust in the flat plate simulation up to Reynolds numbers 10^6 and the best alternative in the driven cavity simulation for Reynolds numbers above 10^3 . The alternating-plane approach exhibits a better behavior for lower Reynolds numbers (below to 10^3) in the driven cavity simulation. A parallel variant of the smoother, tri-plane ordering, presents a good trade-off between convergence and parallel properties.

Key words. plane implicit smoothers, symmetric coupled Gauss-Seidel, robust multigrid, defect correction, Navier-Stokes

Subject classification. Applied and Numerical Mathematics

1. Introduction. Multigrid techniques are generally accepted as fast and efficient methods for solving many types of partial differential equations, and particularly elliptic problems whose discretization results in a K-matrix [24]. For this kind of problems, basic point-wise iterative methods, like Gauss-Seidel or damped Jacobi, are good smoothers, and multigrid methods exhibit an optimal complexity (work is linearly proportional to the number of unknowns), optimal memory requirements, and good parallel efficiency and scalability in parallel implementations [11].

However the efficiency of the multigrid methods degenerates dramatically in presence of anisotropies. It is well known that in the resolution of the Poisson equation the convergence factor of the multigrid method tends to one as the anisotropies are increased [1]. Typically these anisotropies might occur when the coefficients of the discrete operator vary throughout the domain or when stretched grids are used.

*This research was supported by the National Aeronautics and Space Administration under NASA Contract No. NAS1-97046 while the authors were in residence at the Institute for Computer Applications in Science and Engineering (ICASE), NASA Langley Research Center, Hampton, VA 23681-2199.

[†]Departamento de Arquitectura de Computadores y Automática, Universidad Complutense, 28040 Madrid, Spain (email: rubensm@dacya.ucm.es)

[‡]Departamento de Arquitectura de Computadores y Automática, Universidad Complutense, 28040 Madrid, Spain (email: llorente@dacya.ucm.es)

This anisotropic condition occurs naturally in the field of Computational Fluid Dynamics (CFD) where the simulation of small scale physical phenomena, such as boundary layers at high Reynolds numbers, causes the grid to be highly stretched leading to a slow down in convergence.

In some situations, when the source of the anisotropy is known beforehand, a block implicit smoother can be used to improve the efficiency of the multigrid algorithm. Usually this is done by applying a implicit solver in the directions of strong coupling, as states Brandt's fundamental block relaxation rule [1]. This technique is common practice in CFD. Thomas, Diskin and Brandt [20] have demonstrated the efficiency of the distributive smoothing scheme with line solvers applied to high Reynolds number simulations when the grid stretching is normal to the body. The benefits of plane relaxation are shown by Oosterlee in [17] for simulations of the 3-D incompressible Navier-Stokes equations over grids with non-unitary aspect ratios. Also a combination of line implicit techniques and semi-coarsening has been successfully used by Mavriplis in [4] and [5] to solve high Reynolds number 2-D and 3-D viscous flows over anisotropic unstructured meshes.

However, in a general situation the nature of the anisotropy is not known beforehand, so there is no way of knowing which of the variables are coupled. Moreover, if the problem is solved on a stretched grid or the equation coefficients differ from each other throughout the domain (computational and physical anisotropy respectively) the values of the coefficients and their relative magnitudes vary for different parts of the computational domain. In such cases the multigrid techniques based on point- or plane-wise smoothers combined with full coarsening fail to smooth error components with the consequent deterioration of the multigrid convergence factor.

Several approaches aimed at making multigrid a robust solver have been proposed in literature. One popular approach is to use standard coarsening combined with an alternating-direction implicit smoother [19, 10, 12, 19]. This solution consists in exploring all the possibilities in order to develop a robust smoother, i.e. use alternating-line relaxation in 2-D and alternating-plane relaxation in 3-D. Another approach to dealing with anisotropic problems is to combine an implicit smoother with an appropriate semi-coarsening procedure [6, 18]. This is rather popular in literature and overcomes some parallelization problems that can be found in the alternating-plane smoothers [13]. For example, a simple way to avoid using an alternating-plane smoother is to use semi-coarsening in one direction and relaxation in a fixed plane (e.g. combine xy-plane relaxation with Z semi-coarsening). Other intermediate alternatives that combine plane, line or point relaxations with partial and full coarsening have also been presented in multigrid literature [14, 16].

Some of these robust multigrid approaches have also been tested for the efficient resolution of the 2-D Navier-Stokes equations. The alternating-direction line smoother has been investigated for the solution of the incompressible 2-D Navier-Stokes equations in [21, 15]. However, to the authors' knowledge, the robust multigrid algorithms have never been applied to the resolution of the 3-D incompressible Navier-Stokes equations. The aim of this work is to present a thorough study of the application of two common robust multigrid algorithms (alternating-plane smoothers combined with standard coarsening and plane smoothers combined with semi-coarsening) to the resolution of the 3-D Navier-Stokes equations on single-block structured grids.

The robustness of a smoother is defined as its ability to efficiently solve a wide range of problems. In this sense the definition of robustness is qualitative and has to be defined more precisely by setting up a set of suitable test problems. Traditionally, the above mentioned approaches have shown to be robust smoothers for the anisotropic diffusion equation. In the present context we will characterize the multigrid algorithms as robust if the solution of the governing system of equations can be attained in a fixed amount of work units (time to compute the system metrics in the finest level) independent of the grid size, grid stretching factor

and Reynolds number. This is equivalent to saying that the convergence factor of the multigrid algorithm is independent of the grid size, stretching and Reynolds number. We will also refer to this property as *textbook multigrid convergence* (TMC). We believe that the achievement of a textbook multigrid convergence rate through increasing the work and memory requirements per cycle is the first step to achieving *textbook multigrid efficiency* (TME). Opposed to TMC the TME, defined by Brandt in [1], fixes the computational work to solve the problem to ten or less work units.

This paper is organized as follows: The numerical scheme used by our simulations is described in Section 2. Details of the implementation of the multigrid algorithm used in this work will be presented in Section 3. Numerical results are obtained in Section 4 for two common benchmarks in CFD; the driven cavity and the flow over a flat plate. In this section, the robustness of the alternating-plane smoothers combined with full coarsening and plane smoothers combined with semi-coarsening will be investigated. The paper ends with some conclusions in Section 5.

2. The Primitive Equations. The dimensionless steady-state incompressible Navier-Stokes equations in the absence of body forces may be written as:

$$(2.1) \quad \begin{aligned} (\mathbf{u} \cdot \nabla) \mathbf{u} &= -\nabla p + \frac{1}{Re} \Delta \mathbf{u}, \\ \nabla \cdot \mathbf{u} &= 0, \end{aligned}$$

where $\mathbf{u} \in \mathbb{R}^3 = (u, v, w)$ is the non dimensional velocity field and p is the dimensionless pressure. Re is the Reynolds number defined as $Re = \frac{U_\infty \cdot L}{\nu}$, where U_∞ is a characteristic velocity, L a characteristic length and ν the kinematic viscosity.

2.1. Discretization. In order to obtain the discrete expression of the non-linear system (2.1), the solution domain is divided into a finite set of control-volumes (CV). In the present work we will use an orthogonally structured grid where each control volume will be an hexahedron, as in left-hand chart in figure 2.1. The variables are stored in a staggered way, i.e. the velocities are evaluated in the faces of the CV and the pressure field at the center of each CV. Staggered discretization has the benefits of stability properties and leads to a natural discrete form of the continuity equation [7, 9].




FIG. 2.1. Placement of the unknowns in the CV (left-hand chart). Control Volume where the u -momentum equation is integrated (right-hand chart).

The procedure carried out to discretize the u momentum equation will now be described with some detail.

In a staggered arrangement of unknowns each equation is integrated in its own CV. The u -momentum CV is built surrounding the u_{ijk} variable, displaced from the CV of the continuity equation, as in right-hand chart in figure 2.1. In the following, we will refer to the dimensions of this CV Ω as $\Delta X, \Delta Y, \Delta Z$. So we can write the u -momentum equation for a generic node u_{ijk} in its integral form as :

$$(2.2) \quad \int_{\Omega} \mathbf{u} \nabla u \, dV = - \int_{\Omega} \nabla p \, dV + \frac{1}{Re} \int_{\Omega} \Delta u \, dV.$$

The convective term in the momentum equation using the Gauss theorem is rewritten as:

$$(2.3) \quad \int_{\Omega} \mathbf{u} \nabla u \, dV = \int_{\partial\Omega} u (\mathbf{u} \cdot \mathbf{n}) \, dS = \sum_k \int_{\partial\Omega_k} u (\mathbf{u} \cdot \mathbf{n}) \, dS, \quad k = e, w, s, n, t, b.$$

The last integral in (2.3) is easily approximated applying the midpoint rule. Providing the value of the function in the middle of the face results in the leading truncation term being $\mathcal{O}(h^2)$. Thus, to preserve this accuracy the interpolation of the fluxes at the CV faces has to be at least of second-order. This is assured by using a parabolic interpolation for the velocities and linear interpolation for the mass fluxes. Moreover, for non-uniform grids the fluxes are not computed at the middle of the CV. So, assuming a stretched geometric grid of the form $h_{k+1} = \beta h_k$ with β as the grid stretching factor, the integral approximation will have a truncation error $\mathcal{O}((\beta - 1)h) + \mathcal{O}(h^2)$. Taking these considerations into account integral (2.3) is written as:

$$(2.4) \quad \sum_k \int_{\partial\Omega_k} u (\mathbf{u} \cdot \mathbf{n}) \, dS \approx \sum_k m_k u_k, \quad k = e, w, s, n, t, b;$$

where the mass fluxes m_k have been defined as $\int_{\partial\Omega_k} \mathbf{u} \cdot \mathbf{n} \, dS$ and can be evaluated with the following expressions:

$$\begin{aligned} m_e &= \frac{u_{ijk} + u_{i-1jk}}{2} \Delta Y \Delta Z, & m_w &= \frac{u_{ijk} + u_{i+1jk}}{2} \Delta Y \Delta Z, \\ m_s &= \frac{v_{ijk} \Delta x_i + v_{i-1jk} \Delta x_{i+1}}{2} \Delta Z, & m_n &= \frac{v_{ij+1k} \Delta x_i + v_{i-1j+1k} \Delta x_{i+1}}{2} \Delta Z, \\ m_b &= \frac{w_{ijk} \Delta x_i + w_{i-1jk} \Delta x_{i+1}}{2} \Delta Y, & m_t &= \frac{w_{ijk+1} \Delta x_i + w_{i-1jk+1} \Delta x_{i+1}}{2} \Delta Y, \end{aligned}$$

with $\Delta x_{i+1} = x_{i+1jk} - x_{ijk}$ and $\Delta x_i = x_{ijk} - x_{i-1jk}$. The velocity at the CV face is interpolated by fitting a parabola to the values of the velocity at three consecutive nodes: the two nodes located on either side of the surface of interest, plus the adjacent node in the upstream direction. In this work we will use the QUICK formulation of Hayase et al. [8], that can be seen as a defect-correction scheme based on the upwind difference approximation :

$$\begin{aligned} u_e &= \begin{cases} u_{ijk} + S_e^+ (\mathbf{u} \cdot \mathbf{n})_e > 0, \\ u_{i+1jk} + S_e^- (\mathbf{u} \cdot \mathbf{n})_e < 0, \end{cases} & u_w &= \begin{cases} u_{ijk} + S_w^+ (\mathbf{u} \cdot \mathbf{n})_w > 0, \\ u_{i-1jk} + S_w^- (\mathbf{u} \cdot \mathbf{n})_w < 0, \end{cases} & u_n &= \begin{cases} u_{ijk} + S_n^+ (\mathbf{u} \cdot \mathbf{n})_n > 0, \\ u_{ij+1k} + S_n^- (\mathbf{u} \cdot \mathbf{n})_n < 0, \end{cases} \\ u_s &= \begin{cases} u_{ijk} + S_s^+ (\mathbf{u} \cdot \mathbf{n})_s > 0, \\ u_{ij-1k} + S_s^- (\mathbf{u} \cdot \mathbf{n})_s < 0, \end{cases} & u_t &= \begin{cases} u_{ijk} + S_t^+ (\mathbf{u} \cdot \mathbf{n})_t > 0, \\ u_{ijk+1} + S_t^- (\mathbf{u} \cdot \mathbf{n})_t < 0, \end{cases} & u_b &= \begin{cases} u_{ijk} + S_b^+ (\mathbf{u} \cdot \mathbf{n})_b > 0, \\ u_{ijk-1} + S_b^- (\mathbf{u} \cdot \mathbf{n})_b < 0, \end{cases} \end{aligned}$$

The defect-correction source terms S^+ and S^- are calculated within the multigrid cycle using the current approximation whenever a discrete evaluation of the residual is needed. So the algebraic coefficients for the convection terms can be written as:

$$(2.5) \quad \begin{aligned} L_e^c &= \min(0, m_e), & L_w^c &= \min(0, m_w), & L_n^c &= \min(0, m_n), \\ L_s^c &= \min(0, m_s), & L_t^c &= \min(0, m_t), & L_b^c &= \min(0, m_b), \\ L_p^c &= -(L_e^c + L_n^c + L_s^c + L_b^c + L_t^c + L_w^c). \end{aligned}$$

The expression for L_p^c has been obtained using the continuity equation over the CV Ω , which in its discrete form is:

$$m_e + m_w + m_n + m_s + m_t + m_b = 0.$$

Using the Gauss theorem in the diffusive part of the momentum equation (2.3) and the midpoint rule to approximate the resulting surface integral we get:

$$(2.6) \int_{\partial\Omega} \nabla u \cdot \mathbf{n} dS \approx \left[\left(\frac{\partial u}{\partial x} \right)_e - \left(\frac{\partial u}{\partial x} \right)_w \right] \Delta S_x + \left[\left(\frac{\partial u}{\partial y} \right)_n - \left(\frac{\partial u}{\partial y} \right)_s \right] \Delta S_y + \left[\left(\frac{\partial u}{\partial z} \right)_t - \left(\frac{\partial u}{\partial z} \right)_b \right] \Delta S_z.$$

The derivatives in the above expression are evaluated with a central difference scheme :

$$(2.7) \quad \begin{aligned} L_e^d &= \frac{\Delta Y \Delta Z}{Re(x_{i+1jk} - x_{ijk})}, \quad L_w^d = \frac{\Delta Y \Delta Z}{Re(x_{ijk} - x_{i-1jk})}, \quad L_n^d = \frac{\Delta X \Delta Z}{Re(y_{ij+1k} - y_{ij-1k})}, \\ L_s^d &= \frac{\Delta X \Delta Z}{Re(y_{ij+2k} - y_{ijk})}, \quad L_t^d = \frac{\Delta Y \Delta X}{Re(z_{ijk+1} - z_{ijk-1})}, \quad L_b^d = \frac{\Delta Y \Delta X}{Re(z_{ijk+2} - z_{ijk})}, \\ L_p^d &= -(L_e^d + L_n^d + L_s^d + L_b^d + L_t^d + L_w^d). \end{aligned}$$

Finally treating the pressure as a surface force the volume integral in (2.2) can be expressed as a surface integral, as in (2.8). Again this is evaluated using the midpoint rule approximation, in this case no interpolation is needed for the pressure due to the staggered arrangements of unknowns as can be seen in right-hand chart in figure 2.1.

$$(2.8) \quad - \int_{\partial\Omega} p \mathbf{i} \cdot \mathbf{n} dS \approx (p_w - p_e) \Delta S_x \quad \mathbf{i} = (1, 0, 0).$$

Now, we can write the algebraic equation for a generic velocity node u_{ijk} as:

$$(2.9) \quad L_w^u u_{i-1jk} + L_n^u u_{i+1jk} + L_s^u u_{ij-1k} + L_t^u u_{ijk+1} + L_b^u u_{ijk-1} + L_p^u u_{ijk} + L_p^p p_{ijk} + L_w^p p_{i-1jk} = F_{ijk}.$$

The coefficients multiplying the velocity u are obtained as the sum of the diffusive and convective part, i.e. $L_l^u = L_l^c + L_l^d$ with $l = e, w, n, s, b, t$ and those multiplying the pressure are obtained directly from (2.8). An Equivalent expression may be obtained for the v and w momentum equation and can be derived by symmetry from the above equations.

The continuity equation can be easily approximated due to the fact that all velocities are known within the surface of the volume.

$$(2.10) \quad \int_{\Omega} \nabla \cdot \mathbf{u} dV \approx (u_e - u_w) \Delta Y \Delta Z + (v_n - v_s) \Delta X \Delta Z + (w_t - w_b) \Delta X \Delta Y.$$

The above expressions are valid for CV's inside the domain and must be modified in order to satisfy the boundary conditions. The discretization of the boundary conditions are performed by mirroring the cells adjacent to the boundary. The new variables outside the solution domain are extrapolated invoking the boundary condition at each boundary. With these modifications of the algebraic equations the system of non-linear equations to be solved can be presented in a matrix form as:

$$(2.11) \quad \begin{pmatrix} L_u^h & 0 & 0 & L_p^h \\ 0 & L_v^h & 0 & L_p^h \\ 0 & 0 & L_w^h & L_p^h \\ L_m^h & L_m^h & L_m^h & 0 \end{pmatrix} \begin{pmatrix} u \\ v \\ w \\ p \end{pmatrix} = \begin{pmatrix} f_u \\ f_v \\ f_w \\ f_p \end{pmatrix},$$

where the source terms f_u, f_v, f_w and f_p in the right-hand side of the system (2.11) include the discretization of the boundary conditions and the contribution of the QUICK scheme.

3. The Multigrid Method. A sequence of grids $\Omega^l (l = 1, \dots, M)$ is used in the full multigrid (FMG) scheme [1] where Ω^1 is the finest target grid and the rest of the grids are obtained by applying cell-centered coarsening. The computations are initiated in the coarsest grid, once the discrete system is solved the solution is transferred to the next finer level. The prolonged solution is then used as an initial guess for the multigrid method in that level, this procedure is repeated until the finest grid is reached. The goal of this algorithm is to reduce the algebraic error to below the discretization error in just one FMG cycle.

Due to the non-linearity of the problem a Full Approximation Scheme (FAS) [1] is used to solve each level in the FMG cycle. The following iterative algorithm represents a FAS $V(\gamma_1, \gamma_2)$ -cycle to solve the nonlinear system $Lu = f$ on Ω^1 where γ_1 and γ_2 represent the number of pre-smoothing and post-smoothing iterations respectively:

step 1 (*Pre-smoothing*): Apply γ_1 iterations of the smoothing method to $L^1 u^1 = f^1$

FOR $l=1$ TO L *Restriction Part*

step 2: Compute the residual $r^{l-1} = f^{l-1} - L^{l-1} u^{l-1}$

step 3: Restriction of the residual $r^l = R_{l-1}^l r^{l-1}$

step 4: Restriction of the solution $u_{old}^l = I_{l-1}^l u^{l-1}$

step 5: Compute the metrics of level l $L^l(u_{old}^l)$

step 6: Calculate the new right-hand side $f^l = r^l + L^l u_{old}^l$

step 7 (*Pre-smoothing*): Apply γ_1 iterations of the smoothing method to $L^l u^l = f^l$

FOR $l=L-1$ TO 1 *Prolongation Part*

step 8: Correction of the current approximation $u^l = u^l + P_l^{l+1}(u^{l+1} - u_{old}^{l+1})$

step 9: Compute the metrics of level l $L^l(u^l)$

step 10 (*Post-smoothing*): Apply γ_2 iterations of the smoothing method to $L^l u^l = f^l$

In steps 5 and 9 the metrics of the system are computed over the current grid, which includes the computation of the correction terms from the QUICK scheme and also the linearization of the system based on the actual solution. Note that the metrics of the system are also updated within the smoothing process in steps 7 and 10 as explained subsequently in Section 3.1. The operators R_{l-1}^l and P_l^{l+1} in steps 3 and 8 are used to transfer data (solution and residuals) between two different grids; from the coarser level to the current (prolongation) or from the finer to the current level (restriction), respectively.

These transfer operators are dictated by the staggered arrangement of unknowns and the coarsening procedure used. The prolongation and restriction operators are volume-weighted trilinear interpolation in the case of standard coarsening. For the semi-coarsening approach the velocity component parallel to the coarsened direction is restricted using injection, the other variables are restricted using volume-weighted linear interpolation in lines parallel to the coarsened direction. The prolongation operator in this case is volume-weighted linear interpolation. Note that for semi-coarsening the velocity component parallel to the coarsened direction is treated in a vertex-centered way, while the rest of the variables are transferred as cell-centered.

In the following experiments F-cycles will be used (see figure 3.1) to solve each level of the FMG algorithm. F-cycles have been reported to be more efficient for rotating problems [15] at the expense of their parallel properties [4]. The coarsest level is fixed as coarse as possible, and it will be solved with 5 iterations

of the smoothing process.




FIG. 3.1. Scheme of an F-cycle $F(\gamma_1, \gamma_2)$ where γ_0 represents the number of iterations of the smoother performed to solve the coarsest level

3.1. Smoothing Process. One of the most important parts of a multigrid algorithm is the smoothing process. Several smoothers for the Navier-Stokes equations problem were studied in literature. These approaches fall into two categories: (1) coupled smoothing [22, 17, 21] (where the momentum and continuity equation are satisfied simultaneously), and (2) distributive smoothing [20, 3] (where the momentum equations are solved in a first step, and then the velocities and pressures are corrected in order to satisfy the continuity equation). In situations where the coefficients vary through the control volume (e.g. stretched grids, strong recirculating flows,...) coupled smoothing has advantages over the distributive approach because the linearized momentum and continuity equations are solved simultaneously [22, 9]. However the computational cost of the coupled method is much higher than that of the distributive. Note that a (small) matrix has to be inverted in each CV. Moreover, every velocity component is updated essentially twice since it updates all the variables involved in a CV simultaneously (see right-hand chart in figure 2.1).

In particular, we have chosen a cell-implicit Symmetric Coupled Gauss Seidel (SCGS) method as the base smoother because of its higher stability and rapid convergence. This smoother was introduced by Vanka [22] and previously considered in other work [21]. Considering the CV ijk , the momentum equations for the six cell faces together with the continuity equation for the control-volume can be expressed as:

$$\begin{aligned}
 \sum_{|m|+|n|+|p|<1} L_{i+m,j+n,k+p}^{u_w} u_{i+m,j+n,k+p} + L_{p_i}^{u_w} p_{ijk} + L_{p_{i-1}}^{u_w} p_{i-1jk} &= f_{ijk}^u, \\
 \sum_{|m|+|n|+|p|<1} L_{i+m,j+n,k+p}^{u_e} u_{i+m,j+n,k+p} + L_{p_i}^{u_e} p_{i+1jk} + L_{p_{i-1}}^{u_e} p_{ijk} &= f_{i+1jk}^u, \\
 \sum_{|m|+|n|+|p|<1} L_{i+m,j+n,k+p}^{v_s} v_{i+m,j+n,k+p} + L_{p_j}^{v_s} p_{ijk} + L_{p_{j-1}}^{v_s} p_{ij-1k} &= f_{ijk}^v, \\
 \sum_{|m|+|n|+|p|<1} L_{i+m,j+n,k+p}^{v_n} v_{i+m,j+n,k+p} + L_{p_j}^{v_n} p_{ij+1k} + L_{p_{j-1}}^{v_n} p_{ijk} &= f_{ij+1k}^v, \\
 \sum_{|m|+|n|+|p|<1} L_{i+m,j+n,k+p}^{w_b} w_{i+m,j+n,k+p} + L_{p_k}^{w_b} p_{ijk} + L_{p_{k-1}}^{w_b} p_{ijk-1} &= f_{ijk}^w, \\
 \sum_{|m|+|n|+|p|<1} L_{i+m,j+n,k+p}^{w_t} w_{i+m,j+n,k+p} + L_{p_k}^{w_t} p_{ijk+1} + L_{p_{k-1}}^{w_t} p_{ijk} &= f_{ijk+1}^w, \\
 (u_{i+1jk} - u_{ijk})\Delta Y \Delta Z + (v_{ij+1k} - v_{ijk})\Delta X \Delta Z + (w_{ijk+1} - w_{ijk})\Delta X \Delta Y &= f_{ijk}^m.
 \end{aligned}
 \tag{3.1}$$

This set of equations for the CV is linearized by computing the mass fluxes, $L^{u,v,w}$, with the current values of the velocity field. Defining the residuals $r^{u,v,w}$ and the corrections $\Delta u = u^{n+1} - u^n$, etc. the system

(3.1) can be arranged in a block structure as follows:

$$(3.2) \quad \begin{pmatrix} L_{ijk}^{uw} & 0 & 0 & 0 & 0 & 0 & L_{p_i}^{u_w} \\ 0 & L_{i+1,jk}^{ue} & 0 & 0 & 0 & 0 & L_{p_{i-1}}^{u_e} \\ 0 & 0 & L_{ijk}^{vs} & 0 & 0 & 0 & L_{p_j}^{v_s} \\ 0 & 0 & 0 & L_{ij+1,k}^{vn} & 0 & 0 & L_{p_{j-1}}^{v_n} \\ 0 & 0 & 0 & 0 & L_{ijk}^{wb} & 0 & L_{p_k}^{w_b} \\ 0 & 0 & 0 & 0 & 0 & L_{ijk+1}^{wt} & L_{p_{k-1}}^{w_t} \\ -\Delta Y \Delta Z & \Delta Y \Delta Z & -\Delta X \Delta Z & \Delta X \Delta Z & -\Delta X \Delta Y & \Delta X \Delta Y & 0 \end{pmatrix} \begin{pmatrix} \Delta u_{ijk} \\ \Delta u_{i+1,jk} \\ \Delta v_{ijk} \\ \Delta v_{ij+1,k} \\ \Delta w_{ijk} \\ \Delta w_{ijk+1} \\ \Delta p_{ijk} \end{pmatrix} = \begin{pmatrix} r_{ijk}^u \\ r_{i+1,jk}^u \\ r_{ijk}^v \\ r_{ij+1,k}^v \\ r_{ijk}^w \\ r_{ijk+1}^w \\ r_{ijk}^m \end{pmatrix}.$$

A more implicit version of the system (3.2) that includes off-diagonal elements in the first six rows is also possible, this is equivalent to considering implicitly in the equations (3.1) all the unknowns involved in the control-volume. However the convergence factor is similar and the system is more expensive to solve than the system of equations (3.2) [21]. The system (3.2) is easily solved by Gaussian elimination and then the velocity components and the pressure of the CV are updated using under-relaxation:

$$(3.3) \quad \begin{aligned} \mathbf{u}^{n+1} &= \mathbf{u}^n + \omega_u \Delta \mathbf{u}, \\ p^{n+1} &= p^n + \omega_p \Delta p. \end{aligned}$$

The under-relaxation technique has the effect of adding a pseudo-time dependent term in the equations. In the following simulations the under-relaxation factor for the pressure, ω_p , has been fixed to 1.0, while the under-relaxation factor for the velocities, ω_u , is strongly problem dependent and has to be set empirically. The optimum value of ω_u is a function of the Reynolds number, the grid size and also depends on whether the convection scheme is first (upwind) or second order (QUICK) accurate. This is a drawback of this smoother, since a simulation has to be tuned in order to find out the best under-relaxation factor. As figure 3.2 shows the efficiency of the method can be dramatically worsened with a bad choice of ω_u .




FIG. 3.2. Number of fine grid cycles required to converge a driven cavity simulation with $Re=500$ on a $16 \times 16 \times 16$ uniform grid as a function of the relaxation parameter ω_u .

3.2. Plane Implicit Smoothers. Implicit solvers have been widely considered in previous work as a cure to eliminating all the high-frequency errors in the presence of strong anisotropies. Taking advantage of the relatively small 1-D problem size, these implicit line smoothers are based on an exact solver. However the 3-D counterpart does not present this possibility, since the 2-D problem size is no longer small enough to consider using an exact solver. Furthermore a direct exact solver for the planes is not needed, as has

been shown in [10] for the 3-D Poisson equation and in [17] for the incompressible Navier-Stokes equations. This consideration drastically reduces the computational cost of the overall algorithm compared to that of an “exact” plane solver. However this inexact solution of the planes does not decrease the convergence of the multigrid algorithm [17, 10], i.e. solving the plane beyond a precision threshold does not improve the convergence rate. Note that the plane implicit smoother has to damp high oscillating error components in the plane rather than solving a 2-D problem exactly.

In the present work, the planes will be approximately solved with a 2-D multigrid algorithm consisting of one FAS F(1,1) cycle. The same kind of anisotropies found in the 3-D problem may appear in the 2-D system. Thus a robust multigrid algorithm is, again, completely necessary. For the 2-D system the same robust algorithms will be considered, i.e. an alternating-line smoother combined with full coarsening and a line-implicit smoother combined with semi-coarsening. One 1-D FAS F(1,1) cycle will be applied to solve the lines, the smoother in this case being SCGS as described in the previous section.

The coupled philosophy of the SCGS will be applied in the line and plane solvers. The plane smoother relaxes simultaneously the momentum and continuity equations of the cells included in the plane, and hence all velocity components and pressures contained within the plane will be updated at the same time. Let us consider for example an xy plane, defining the vector X_k that accommodates the variables for a whole plane of cells:

$$X_k^T = (\mathbf{u}, \mathbf{v}, \mathbf{w}, \mathbf{w}^+, \mathbf{p}), \mathbf{u} = u_{ijk}, \mathbf{v} = v_{ijk}, \mathbf{w} = w_{ijk}, \mathbf{w}^+ = w_{ijk+1}, \mathbf{p} = p_{ijk}, \quad \forall i, j \in [0, n] \quad k = \text{const.}$$

The equation system for the plane in terms of residuals and corrections is:

$$(3.4) \quad L_k \Delta X_k = R_k,$$

where $R_k = f_k - L_k X_k$ is the residual of the k^{th} plane and $\Delta X_k = X_k^{n+1} - X_k^n$ is the increment of the solution. The system of equations (3.4) is built into the smoothing process as follows. When solving the k^{th} plane, the metrics in that plane are linearized using the current solution. Also the second-order correction for the convective term is recomputed. With these new metrics the residual R_k for the k^{th} plane is calculated.

Defining a specific ordering of the planes, many types of plane smoothers can be easily constructed. It is important to note that with the second order operator, the right-hand side of the system (3.4) depends on the values of the plane k , $k \pm 1$ and/or $k \pm 2$ depending on the direction of the velocity. Thus a parallel implementation can not be constructed based on a regular zebra ordering (left-hand chart in figure 3.3). In order to avoid these dependencies a tri-plane smoother could be applied [15] (right-hand chart in figure 3.3)

Grid stretching is commonly used in grid generation to pack points into regions with large solution gradients while avoiding an excess of points in more benign regions (for example in the simulation of viscous flows at high Reynolds number to resolve boundary layers). The convergence of multigrid based on point smoothing and full coarsening deteriorates dramatically when highly stretched grids are used. In some situations, when the direction of the anisotropies is known beforehand, the multigrid convergence can be improved using an implicit smoother in the direction normal to the stretching. However, if the stretched grid generates aspect ratios whose relative magnitudes vary for different parts of the computational domain the multigrid techniques based on plane-wise smoothers combined with full coarsening fail to smooth error components. Other remedies should be used to achieve a robust solver, the two most common alternatives being:

- *Robust multigrid smoothing process with standard coarsening.* If the coarser grids are built by doubling the mesh size in all coordinates direction, sweeps of the planes in the three directions are

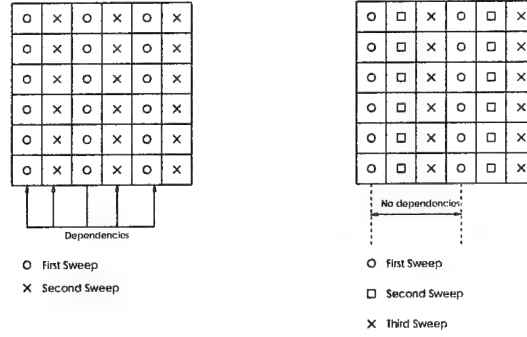


FIG. 3.3. Standard zebra ordering of planes (left-hand chart). Planes relaxed concurrently in a tri-plane smoother (Right-hand chart).

needed to achieve robustness ((y,z)-plane smoothing sweep \rightarrow (x,z)-plane smoothing sweep \rightarrow (x,y)-plane smoothing sweep). From here on, this approach will be referred as alternating-plane smoothers (APS). Several versions of this method can be developed depending on the sweep ordering: symmetric alternating-plane smoother (S-APS), lexicographic alternating-plane smoother (L-APS) and tri-plane alternating-plane smoother (Tri-APS).

- *Plane implicit smoothers combined with semi-coarsening.* Instead of using standard coarsening the coarse levels can be built by only coarsening along one direction. In order to achieve robustness a plane implicit solver perpendicular to the coarsened direction is needed. Based on the coarsened direction, we will refer to these approaches as X, Y or Z semi-coarsening (XSC, YSC, ZSC). Depending on the order in which the planes are swept we can construct the following methods: symmetric Z semi-coarsening (S-ZSC), lexicographic Z semi-coarsening (L-ZSC) and tri-plane Z semi-coarsening (Tri-ZSC).

4. Numerical Experiments. Two different flows have been chosen to test the robustness of the multigrid algorithms described previously: the driven cavity and the flow over a flat plate. These two cases have been widely studied and used as benchmarking problems for CFD codes. Although the flow structure are relatively simple, they exhibit some basic problems that prevent optimal multigrid efficiencies from being achieved [2], namely strong recirculating flows and boundary layers.

Let R be the L_2 -norm of the average residual of the system of equations (2.11) defined as:

$$(4.1) \quad R = \sqrt{\frac{\sum ((R_{ijk}^u)^2 + (R_{ijk}^v)^2 + (R_{ijk}^w)^2 + (R_{ijk}^c)^2)}{4 \cdot N_x \cdot N_y \cdot N_z}},$$

where R^u , R^v , R^w , R^c are the residuals of the u , v , w momentum equations and continuity equation respectively. The convergence criterion is based on R . When the fine grid average residual decreases to below 10^{-4} the calculations are terminated. This value is small enough to assure that the algebraic error is below to the discretization error. Let R_0 and R_n denote respectively the residual norms (as defined in 4.1) before the iterative process and after the convergence criterion is satisfied. So the average convergence factor is defined by:

$$(4.2) \quad \bar{\rho} = \left(\frac{R_n}{R_0} \right)^{\frac{1}{n}}.$$

4.1. Flow in a Driven Cavity. The numerical solution, which has been widely used for testing numerical schemes, is that of a flow confined in a rectangular domain with the upper wall moving at a

constant speed. The flow structure for *low to moderate* Reynolds numbers consists of a 3-D primary vortex and two 3-D secondary vortices at the bottom. The problem currently considered consists of a cube of dimension L with the top wall moving at a velocity u . The Reynolds number of the flow based in these quantities is $Re = \frac{uL}{\nu}$. The boundary conditions are of Dirichlet type for the velocities on the six faces of the computational domain, and no boundary conditions were necessary for the pressure.




FIG. 4.1. Grid used for the driven-cavity problem 32x32x32 (left-hand chart), and structure of the 3D primary vortex for $Re=10^3$ (right-hand chart).

Simulations have been performed over three different grids, each one uniform and stretched : 16x16x16, 32x32x32 and 64x64x64. The stretched grids were of the form $h_{k+1} = \beta h_k$, the stretching factor being β equal to 1.1 in all cases (see figure 4.1 left-hand chart). The driven cavity problem is a rotating flow for which standard multigrid schemes might have difficulties to converge. These difficulties were not experienced in this work since a moderate Reynolds number was considered. However, the simulations result in a complex re-circulating flow consisting of 3-D vortex structures as can be seen in figure 4.1. The flow field is in good qualitative agreement with previous flow calculations [23]. The profiles at the center line of the cavity with $Re=10^3$ on a 32x32x32 stretched grid can be seen in figure 4.2.




FIG. 4.2. The u velocity component profile at the center line of the cavity with $Re=10^3$ for first-order and second-order accurate solutions over a 32x32x32 stretched grid.

As mentioned in section 3.1 the convergence factor varies with the choice of the under-relaxation factor for the velocity field. In table 4.1 the under-relaxation factors used in each simulation, as a function of the Reynolds number and grid size, are shown for the different multigrid cycle under study: symmetric SCGS,

tri-plane and symmetric APS, and tri-plane and symmetric ZSC.

TABLE 4.1

Under-relaxation factors for the driven-cavity simulation as a function of the grid size, Reynolds number and multigrid cycle.

Reynolds Number	Grid								
	16x16x16			32x32x32			64x64x64		
	SCGS	APS	ZSC	SCGS	APS	ZSC	SCGS	APS	ZSC
10^2	0.6	0.6	0.5	0.5	0.6	0.6	0.5	0.6	0.5
10^3	0.4	0.3	0.4	0.3	0.3	0.3	0.3	0.3	0.3

Figure 4.3 shows the L_2 -norm of the residual versus $F(1,1)$ -cycles with a symmetric SCGS smoother for several uniform grids and Reynolds numbers. The behavior of this smoother is quite good for low Reynolds numbers (left-hand graph in figure 4.3 for Reynolds number 10^2), the residual norm is reduced by between four and five orders of magnitude in the first five cycles. However its efficiency decreases as the problem becomes more convective. The residual norm can not be reduced by four orders of magnitude in ten cycles (right-hand graph in figure 4.3 for Reynolds number 10^3). Furthermore, convergence could not be attained over stretched grids with a cell-wise smoother like symmetric SCGS.




FIG. 4.3. L_2 -norm of the residual versus $F(1,1)$ -cycles with symmetric SCGS smoother for several uniform grids and Reynolds numbers for the driven-cavity simulation.

Figure 4.4 shows the L_2 -norm of the residual versus $F(1,1)$ -cycles with an alternating-plane smoother combined with full coarsening (APS) for several grids and Reynolds numbers. The symmetric APS approach (top graphs in figure 4.4) converges the residual to below 10^{-4} in five cycles for both Reynolds numbers (10^2 and 10^3). However, the tri-plane APS approach (bottom graphs in figure 4.4) need eight cycles to reduce the residual norm to below 10^{-4} for Reynolds number 10^3 . It is interesting to note that the cost per multigrid cycle with the symmetric ordering is about twice as large as that with the tri-plane ordering. However convergence factor per work unit is better with the symmetric ordering of planes for low Reynolds numbers and it is similar for both smoothers for $Re = 10^3$.

One of the drawbacks of the APS approach is its difficult implementation in a parallel setting [13]. This problem can be easily overcome using a plane smoother combined with semi-coarsening to ensure robustness. The block implicit smoother used to converge the driven cavity simulation needs to be applied along the sub-characteristics of the discrete operator for convection dominated problems in order to obtain the higher efficiency. For example, it was observed that the xz -plane sweeps seriously harm the smoothing, and so the

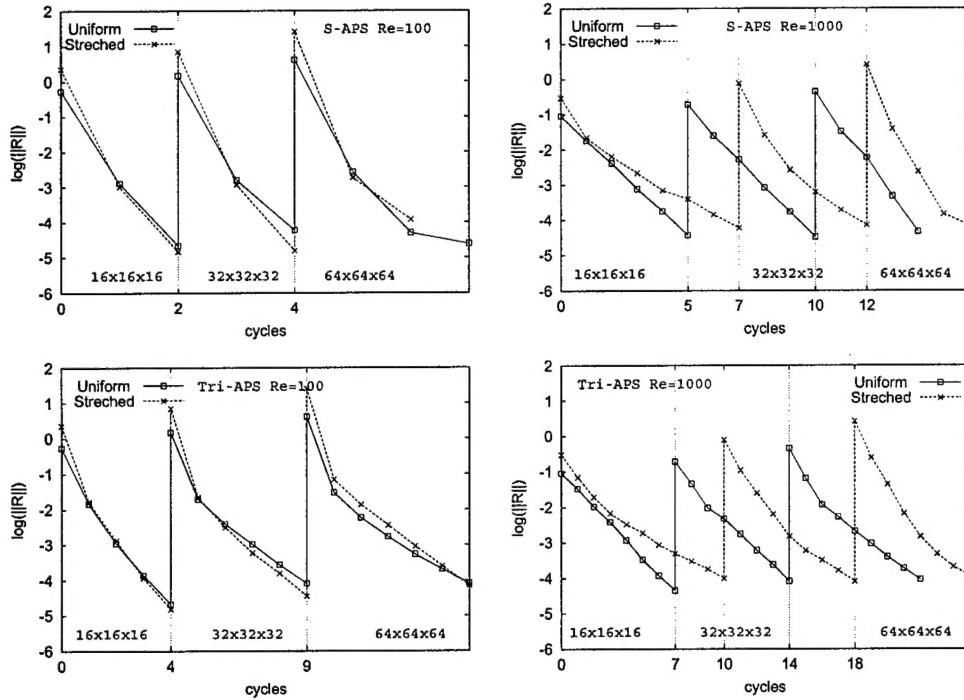


FIG. 4.4. L_2 -norm of the residual versus $F(1,1)$ -cycles with an alternating-plane smoother combined with full coarsening (APS) for several grids and Reynolds numbers with a symmetric order of planes (top graphs) and a tri-plane ordering (low graphs) for the driven cavity simulation.

Y semi-coarsening approach exhibits a poor behavior.

Figure 4.5 shows the L_2 -norm of the residual versus $F(1,1)$ -cycles with an xy-plane implicit smoother combined with Z semi-coarsening (ZSC) for several grids and Reynolds numbers. Although not shown, the behavior exhibited by the XSC approach is similar to the one presented in the set of figures 4.5 for ZSC. The symmetric ZSC approach (top graphs in figure 4.5) converges the residuals to below 10^{-4} in five cycles for both Reynolds numbers (10^2 and 10^3). However, the tri-plane ZSC approach (bottom-charts in figure 4.5) is not able to reduce the residual norm below 10^{-4} in ten cycles for Reynolds number 10^3 . As one might expect, the time per cycle is twice as fast for the tri-plane ordering, although the convergence factor per work unit is better with the symmetric ordering for all the cases.

Table 4.2 shows the average convergence factors obtained in the simulation of the driven cavity problem for several uniform and stretched grids, Reynolds numbers and $F(1,1)$ cycles. The convergence factor has been proved to be independent of the grid size and stretching for the two robust approaches investigated, however the convergence is not Reynolds number independent for the driven cavity simulation. The algorithms exhibit the same behavior for low Reynolds numbers as when solving the Poisson equation, i.e. the residual reduction per cycle is similar in both situations [10]. The convergence factor improves on stretched grids (as in the fully elliptic case [10]), and as shown in table 4.2, it also improves for finer grids. The convergence factor for the APS approach is lower than for the SC approach and its cost per cycle is twice as low because the F-cycle spends a lot of time on coarser levels. However its difficult and low-efficiency parallelization and its difficulty to converge for Reynolds numbers higher than 10^3 might make the semi-coarsening approach more attractive.

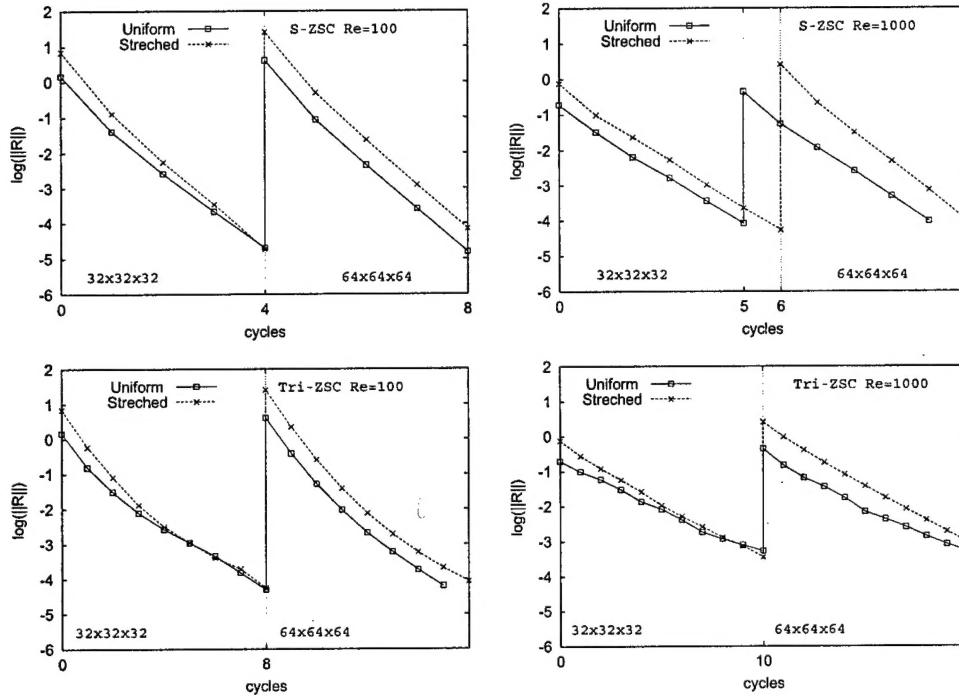


FIG. 4.5. L_2 -norm of the average residual versus $F(1,1)$ -cycles with an xy -plane implicit smoother combined with z semi-coarsening (ZSC) for several grids and Reynolds numbers with a symmetric order of planes (top graphs) and a tri-plane ordering (low graphs) for the driven cavity simulation.

TABLE 4.2

Average convergence factors obtained in the simulation of the driven cavity problem for several uniform (U) and stretched (S) grids, Reynolds numbers and different $F(1,1)$ cycles.

	Grid							
	32x32x32				64x64x64			
	S-APS	S-ZSC	Tri-APS	Tri-ZSC	S-APS	S-ZSC	Tri-APS	Tri-ZSC
$Re = 10^2$ U	$1.6 \cdot 10^{-3}$	0.07	0.14	0.24	$3 \cdot 10^{-3}$	0.04	0.17	0.21
$Re = 10^2$ S	$1.4 \cdot 10^{-3}$	0.04	0.08	0.23	$7 \cdot 10^{-3}$	0.04	0.07	0.20
$Re = 10^3$ U	0.17	0.21	0.25	0.5	0.10	0.18	0.34	0.5
$Re = 10^3$ S	0.15	0.21	0.3	0.46	0.10	0.15	0.34	0.46

4.2. 3D Flat Plate Boundary Layer. We consider an square plate placed in the middle of the solution domain. In the west face ($x = 0$) we define the inflow boundary with no angle of attack, and so the east face will hold the outflow condition. On the plate a no-slip boundary condition is imposed, and symmetric condition is imposed elsewhere on the domain boundary. As the velocity gradient normal to the wall is very high only in the boundary layer, the thin-layer approximation which only retains those terms can be adopted. However in the following simulations the original form (2.1) of the Navier-Stokes equations is solved.

In order to capture the viscous effects, the grid is highly stretched near the plate (see left-hand chart in figure 4.6). Moreover, the grid is refined near the plate edges to reduce the large discretization errors

in those zones as advocated by Thomas et al. [20]. To ensure that a sufficient number of grid points will lie inside the boundary-layer, the mesh space for a uniform mesh would impose too high a demand on the computation. For example, approximating the boundary-layer thickness with $\delta \sim \frac{1}{\sqrt{Re}}$ for $Re = 10^4$ we have $\delta \sim 0.01$ which implies at least 10^2 grid points in a uniform grid which cannot be considered due to memory limitations in a 3-D simulation. Thus for this model problem, no regular grids will be considered. The grids are stretched in the z -direction using a geometric factor $h_k = \beta h_{k-1}$ with $\beta = 1.3$ for the $24 \times 24 \times 32$ grid and $\beta = 1.1$ for the $48 \times 48 \times 64$ grid.




FIG. 4.6. $48 \times 48 \times 32$ grid used for the flat-plate simulation (left-hand chart). Pressure contour lines and boundary layer for $Re = 10^4$ (right-hand chart).

The solution is verified by comparing the u -velocity in the middle of the plate with the Blasius analytical solution for a 2-D plate (figure 4.7). The little discrepancy near the layer edge is due to the highly stretched grid used in this simulation. The second-order accuracy has been verified using the solution in the three finest grids.




FIG. 4.7. Simulation comparison with Blasius theory at the middle of the plate with $Re = 10^4$.

The multigrid cycle employed to solve each level of the FMG is a F(2,1) cycle. The under-relaxation factors used in the simulations are shown in table 4.3. Depending on the problem, some plane sweep directions may deteriorate the smoothing. The best smoothing rate was achieved with a combination of xy-plane relaxation and z semi-coarsening (ZSC). We do not include results of the alternating-plane approach or other semi-coarsening directions because of their poor behavior.

TABLE 4.3

Under-relaxation factors for the flat-plate simulation as a function of the grid size, Reynolds number and multigrid cycle.

Reynolds Number	Grid			
	24x24x32		48x48x64	
	L-ZSC	Tri-ZSC	L-ZSC	Tri-ZSC
10^2	0.8	0.8	0.8	0.8
10^4	0.6	0.4	0.6	0.4

Figure 4.8 shows the L_2 -norm of the residual versus F(2,1)-cycles with lexicographic and tri-plane xy-plane implicit smoothers combined with Z semi-coarsening (ZSC) for several grids and Reynolds numbers. The residual norm is reduced by nearly five orders of magnitude in the first five cycles in all cases (note that the reduction is of four orders of magnitude in the first two cycles for the 48x48x64 grid). In fact, the full multigrid algorithm converges the solution to below the truncation error with one F(2,1) cycle per level. The asymptotic convergence rate is equal to 0.19 which is close to that obtained for the Poisson equation with the semi-coarsened approach [13].

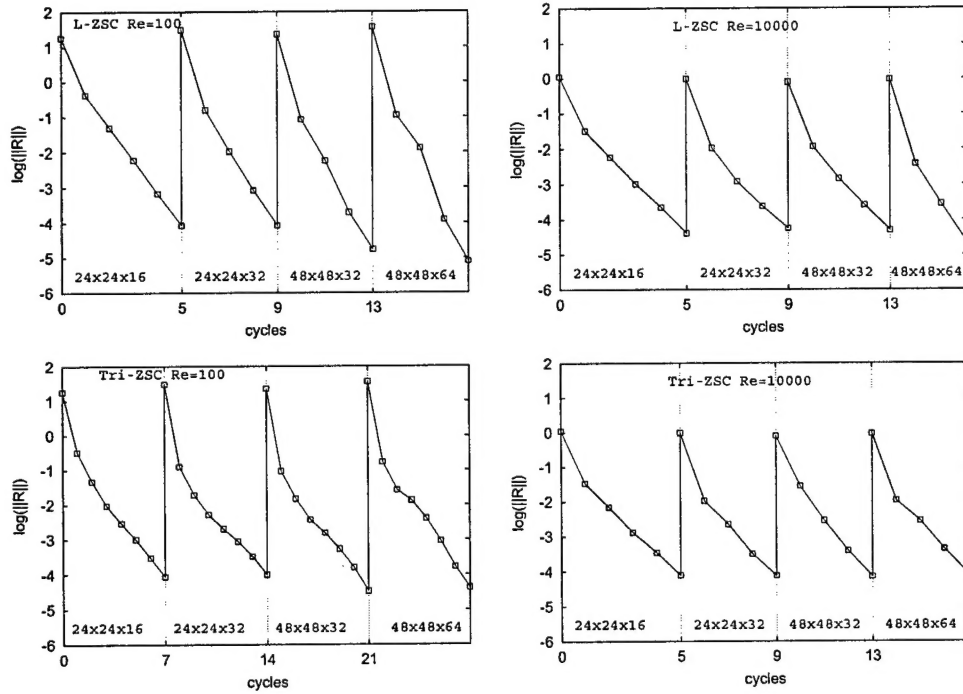


FIG. 4.8. L_2 -norm of the residual versus F(2,1)-cycles with lexicographic (top graphs) and tri-plane (bottom graphs) xy-plane implicit smoothers combined with z semi-coarsening (ZSC) for several grids and Reynolds numbers.

Table 4.4 shows the average convergence factors obtained in the simulation of the flat plate problem for different F(2,1) cycles and several stretched grids and Reynolds numbers. Although not included in this report, experiments with Reynolds numbers up to 10^6 were performed. Convergence rates, independent of the Reynolds number, the grid size and the stretching factor, were achieved for the resolution of the boundary layer over a flat plate. Since the flow is aligned with the grid, the results obtained with the lexicographic